

Web-based Supplemental Material for A Hybrid Phase I-II/III Clinical Trial Design Allowing Dose Re-Optimization in Phase III

Andrew G. Chapple and Peter F. Thall

May 21, 2018

A Review of Eff-Tox Tradeoff Based Phase I-II Design

The phase I-II design introduced by Thall and Cook (2004), and refined by Thall, et al. (2014), may be summarized as follows. All necessary computations for implementation are carried out by the *EffTox* program, available on the MDAnderson biostatistics software website. Indexing efficacy by E and toxicity by T , the bivariate binary outcome, (Y_E, Y_T) , takes on four possible values, (a, b) , with $a, b \in \{0, 1\}$. For each marginal probability, $m = E, T$, and standardized dose x_j , denoting the model parameter vector by $\boldsymbol{\theta}_{ET}$, it is assumed that $\pi_m(x_j, \boldsymbol{\theta}_{ET}) = P(Y_m = 1 \mid x_j, \boldsymbol{\theta}_{ET}) = \text{logit}^{-1}\{\eta_m(x_j, \boldsymbol{\theta}_{ET})\}$, with linear terms

$$\eta_T(x_j, \boldsymbol{\theta}_{ET}) = \tau_{T,1} + \tau_{T,2}x_j \quad \text{and} \quad \eta_E(x_j, \boldsymbol{\theta}_{ET}) = \tau_{E,1} + \tau_{E,2}x_j + \tau_{E,3}x_j^2. \quad (1)$$

with $\tau_{T,2} > 0$. These linear terms are formulated to ensure that $\pi_T(x_j, \boldsymbol{\theta}_{ET})$ increases with dose, but to allow $\pi_E(x_j, \boldsymbol{\theta}_{ET})$ to possibly be non-monotone in dose. An association parameter ψ is used to define the joint distribution of Y_E and Y_T from their marginals using a copula, as follows. Suppressing the arguments $x_j, \boldsymbol{\theta}_{ET}$ for brevity, the four joint probabilities $\pi_{a,b}$ for $a, b \in \{0, 1\}$ are given by

$$\pi_{a,b} = \pi_E^a(1 - \pi_E)^{1-a}\pi_T^b(1 - \pi_T)^{1-b} + (-1)^{a+b}\pi_E(1 - \pi_E)\pi_T(1 - \pi_T) \left(\frac{e^\psi - 1}{e^\psi + 1} \right). \quad (2)$$

Thus $\boldsymbol{\theta}_{ET} = (\tau_{T,1}, \tau_{T,2}, \tau_{E,1}, \tau_{E,2}, \tau_{E,3}, \psi)$. For priors, it is assumed that these six parameters are independent with $\psi \sim N(0, 1)$, $\tau_{E,2} \sim N(0, .20)$, and $\tau_{m,r} \sim N(\tilde{\mu}_{m,r}, \tilde{\sigma}_{m,r}^2)$ for $m = E, T$ and $r = 1, 2$. This requires specifying numerical values of the four hyper-means and four hyper-standard deviations. This is done using elicited means of $\pi_m(x_j, \boldsymbol{\theta}_{ET})$, for $j = 1, \dots, J$, $m = E, T$, with a specified prior effective sample size used to calibrate the priors of the eight remaining hyper-parameters. Details are provided by Thall et al.(2014).

During the trial, a dose x_j is acceptable in phase I-II if

$$P\{\pi_E(x_j, \boldsymbol{\theta}_{ET}) > \underline{\pi}_E \mid \text{data}\} > p_E \quad \text{and} \quad P\{\pi_T(x_j, \boldsymbol{\theta}_{ET}) < \bar{\pi}_T \mid \text{data}\} > p_T,$$

where $\underline{\pi}_E, \bar{\pi}_T, p_E, p_T$ are design parameters fixed in advance.

Adaptive dose-finding decisions are based on a real-valued desirability (trade-off) function $\phi(\pi_E, \pi_T)$ for $\boldsymbol{\pi} = (\pi_E, \pi_T) \in [0, 1]^2$. Key properties of $\phi(\pi_E, \pi_T)$ are that (1) it must increase in π_E , (2) it must decrease in π_T , and (3) it should quantify the desirability of each $\boldsymbol{\pi} \in [0, 1]^2$ as the risk-benefit trade-off between allowing a higher π_T and achieving a higher π_E . To compute ϕ , one may first elicit three equally desirable outcome probability pairs $\boldsymbol{\pi}_1^* = (\pi_{1,E}^*, 0)$, $\boldsymbol{\pi}_2^* = (1, \pi_{2,T}^*)$ and $\boldsymbol{\pi}_3^* = (\pi_{3,E}^*, \pi_{3,T}^*)$, with $\pi_{1,E}^* < \pi_{3,E}^*$ and $\pi_{3,T}^* < \pi_{2,T}^*$. Since $\boldsymbol{\pi} = (1, 0)$ is the optimal probability pair, the function ϕ then may be defined by the equation

$$\phi(\pi_E, \pi_T) = 1 - \|(\pi_E, \pi_T) - (1, 0)\|_p = 1 - \left\{ \left(\frac{\pi_E - 1}{\pi_{E,1}^* - 1} \right)^p + \left(\frac{\pi_T - 0}{\pi_{T,2}^* - 0} \right)^p \right\}^{1/p} \quad (3)$$

where $p > 0$. The equation $\phi(\pi_{E,3}^*, \pi_{T,3}^*) = 0$ is solved for p using the bisection method, to give $\phi(\boldsymbol{\pi}) = 0$ on the target contour \mathcal{C}_0 . This in turn defines a family of trade-off contours in $[0, 1]^2$ such that

$$\mathcal{C}_z = \{(\pi_E, \pi_T) : \phi(\pi_E, \pi_T) = z\},$$

so for real z all pairs on \mathcal{C}_z have the same desirability z . To use this structure for dose selection, denote the marginal hyperparameter vectors by ζ_E and ζ_T with $\mu_{m,j,n}(\zeta_m) = E\{\pi_m(x_j, \boldsymbol{\theta}_m) \mid \mathcal{D}_n, \zeta_m\}$ for $m = E, T$ and x_j . The acceptable dose with largest estimated desirability $\phi(\mu_{E,j,n}(\zeta_E), \mu_{T,j,n}(\zeta_T))$ is chosen. As explained by Thall, et al. (2014), the three probability pairs used to define \mathcal{C}_0 must be chosen to give a family of target contours sufficiently steep so that there is a reasonable payoff for increasing dose, and thus obtaining larger $\mu_{E,k,n}(\tilde{\boldsymbol{\theta}}_E)$ as a trade-off for larger $\mu_{T,k,n}(\tilde{\boldsymbol{\theta}}_T)$. Web Figure 1 illustrates the trade-off contour used in the illustrative trial in the simulation study.

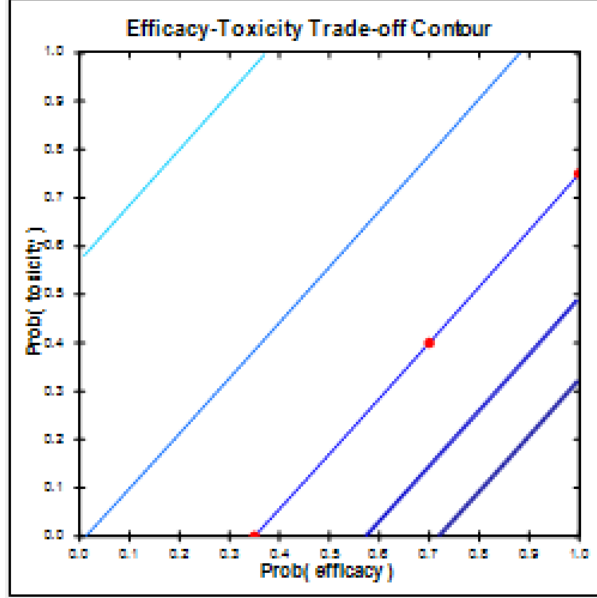


Figure 1: Contour used for determining values of $\phi(\cdot)$ in the simulation study.

B Computational Algorithm for Simulating Posterior Samples

We perform a Markov chain Monte Carlo (MCMC) sampling scheme for $\boldsymbol{\theta}_{ET}$ and $\boldsymbol{\theta}_S$ separately, where the MCMC for $\boldsymbol{\theta}_S$ also includes reversible jump moves. The sampling for $\boldsymbol{\theta}_{ET} = (\tau_{T,1}, \tau_{T,2}, \tau_{E,1}, \tau_{E,2}, \tau_{E,3}, \psi)$ involves Metropolis Hastings moves for each entry with adaptive proposal variances for the first half of the posterior sampling to improve convergence rates. Since we impose the assumption that toxicity probability increases with dose, $\log(\tau_{T,2}^*)$ is proposed from a normal distribution with mean $\log(\tau_{T,2})$ and variance $\sigma_{\tau_{T,2}}$, denoted by $N(\log(\tau_{T,2}), \sigma_{\tau_{T,2}})$, which is adjusted every 250 iterations by either doubling or halving its value during the first half of the MCMC to maintain acceptance rates between .2 and .8. Let \mathbf{Y}_S denote all (Y^0, δ) data. Denoting likelihood by L , the proposed value $\tau_{T,2}^*$ is accepted over the previous value $\tau_{T,2}$ with probability

$$\min \left\{ \frac{\tau_{T,2}^* L(\boldsymbol{\theta}_{ET}^* | \mathbf{Y}_E, \mathbf{Y}_T, \mathbf{x}) N(\tau_{T,2}^* | 0, .2)}{\tau_{T,2} L(\boldsymbol{\theta}_{ET} | \mathbf{Y}_E, \mathbf{Y}_T, \mathbf{x}) N(\tau_{T,2} | 0, .2)}, 1 \right\}.$$

Let $\theta_{ET,j}$ denote an entry of $\boldsymbol{\theta}_{ET}$ for $j = 2, \dots, 6$. We propose a value $\theta_{ET,j}^*$ by sampling from a $N(\theta_{ET,j}, \sigma_{\theta_{ET,j}}^2)$, which is adjusted every 250 iterations by either doubling or halving its value during the first half of the MCMC to maintain acceptance rates between .2 and .8. The proposed value which is adjusted every 250 iterations by either doubling or halving its value during the first half of the MCMC to maintain acceptance rates between .2 and .8. The proposed value $\theta_{ET,j}^*$ and vector $\boldsymbol{\theta}_{ET}^*$ is accepted over the previous value with probability

$$\min \left\{ \frac{L(\boldsymbol{\theta}_{ET}^* | \mathbf{Y}_E, \mathbf{Y}_T, \mathbf{x}) N(\theta_{ET,j}^* | \tilde{\mu}_j, \tilde{\sigma}_j^2)}{L(\boldsymbol{\theta}_{ET} | \mathbf{Y}_E, \mathbf{Y}_T, \mathbf{x}) N(\theta_{ET,j} | \tilde{\mu}_j, \tilde{\sigma}_j^2)}, 1 \right\}$$

where $\tilde{\mu}_j$ is derived from clinician elicited prior expected efficacy and toxicity probabilities at each dose and $\tilde{\sigma}_j^2$ is calibrated so that the design has a desired prior effective sample size (ESS), with a suggested value of ESS near .9. For $\theta_{ET,6} = \psi$, $\tilde{\mu}_6 = 0$ and $\tilde{\sigma}_j^2 = 1$. The other prior means and variances can be determined using the *EffTox* software available on the MD Anderson Biostatistics website. We found that drawing 2000 samples and discarding the first half of the samples, with proposal variances adjusted throughout, showed good convergence as evidenced by parameter traceplots.

In our MCMC sampling scheme for θ_S , we perform five different moves consisting of four Metropolis-Hastings moves (one of which is Metropolis-Hastings-Green move) and one Gibbs sampler. We perform 2000 iterations and burn in the first half, adaptively adjusting proposal variances throughout the first half of the MCMC. Let β_S denote the vector $(\beta_1, \beta_2, \beta_E, \beta_T)$. One generic iteration of the sampler proceeds as follows:

1. Sample $\beta_S | \mathbf{Y}_S, \mathbf{Y}_E, \mathbf{Y}_T, \mathbf{x}, \mathbf{t}, \boldsymbol{\lambda}$ via a Metropolis-Hastings step. For entry j , we propose $\beta_{S,j}^*$ from a $N(\beta_{S,j}, \sigma_{\beta_{S,j}}^2)$, where $\sigma_{\beta_{S,j}}^2$ is halved or doubled every 250 iterations during the burn-in period to ensure that these parameters exhibit good convergence. We accept $\beta_{S,j}^*$ over the previous value with probability equal to the likelihood ratio times $\min\{N(\beta_{S,j}^* | 0, 100) / N(\beta_{S,j} | 0, 100), 1\}$. As before, here the variance is set to 100 to ensure that the prior does not control the decision making.
2. Sample $\sigma^{-2} | \mu, \boldsymbol{\lambda}, \mathbf{t}$ via a Gibbs step. Sample directly from:

$$\sigma_{\lambda}^{-2} | \boldsymbol{\lambda} \sim \text{Gamma} \left(.5J + 1, 0.5 \sum_{j=2}^{J+1} (\lambda_j - \lambda_{j-1})^2 \right)$$

3. Sample $\boldsymbol{\lambda} | \sigma_{\lambda}, \mathbf{Y}_S, \mathbf{Y}_E, \mathbf{Y}_T, \mathbf{x}, \mathbf{t}$ via a Metropolis-Hastings step. We sample the l th entry of $\boldsymbol{\lambda}$ for $l = 1, \dots, J + 1$ by drawing $\lambda_l^* \sim U[\lambda_l - .25, \lambda_l + .25]$ and accept it with probability

$$\frac{L(\boldsymbol{\lambda}^* | \mathbf{Y}_S, \mathbf{Y}_E, \mathbf{Y}_T, \boldsymbol{\beta}_S, \mathbf{x}, \mathbf{t}) \prod_{l=2}^{L+1} N(\lambda_l^* | \lambda_{l-1}, \sigma_{\lambda}^2)}{L(\boldsymbol{\lambda} | \mathbf{Y}_S, \mathbf{Y}_E, \mathbf{Y}_T, \boldsymbol{\beta}_S, \mathbf{x}, \mathbf{t}) \prod_{l=2}^{L+1} N(\lambda_l | \lambda_{l-1}, \sigma_{\lambda}^2)}$$

for $l = 2, \dots, L + 1$ when $L > 0$. We accept moves on λ_1 with probability

$$\min \left\{ \frac{L(\boldsymbol{\lambda}^* | \mathbf{Y}_S, \mathbf{Y}_E, \mathbf{Y}_T, \boldsymbol{\beta}_S, \mathbf{x}, \mathbf{t}) N(\lambda_1^* | 0, 25)}{L(\boldsymbol{\lambda} | \mathbf{Y}_S, \mathbf{Y}_E, \mathbf{Y}_T, \boldsymbol{\beta}_S, \mathbf{x}, \mathbf{t}) N(\lambda_1 | 0, 25)}, 1 \right\} \quad \text{if } L = 0,$$

else

$$\min \left\{ \frac{L(\boldsymbol{\lambda}^* | \mathbf{Y}_S, \mathbf{Y}_E, \mathbf{Y}_T, \boldsymbol{\beta}_S, \mathbf{x}, \mathbf{t}) N(\lambda_1^* | 0, 25) N(\lambda_2 | \lambda_1^*, \sigma_{\lambda}^2)}{L(\boldsymbol{\lambda} | \mathbf{Y}_S, \mathbf{Y}_E, \mathbf{Y}_T, \boldsymbol{\beta}_S, \mathbf{x}, \mathbf{t}) N(\lambda_1 | 0, 25) N(\lambda_2 | \lambda_1, \sigma_{\lambda}^2)}, 1 \right\}$$

The prior variance on λ_1 is set to 25 to be non-informative for the hazard $\exp(\lambda_1)$ on the interval $[0, t_1)$ while also imposing realistic hazard values *a priori*.

4. Sample the locations of $\mathbf{t} | \boldsymbol{\lambda}, \mathbf{Y}_S, \mathbf{Y}_E, \mathbf{Y}_T, L$ via a Metropolis-Hastings move that shuffles the locations of t_1, \dots, t_L . We sample $t_l^* \sim U[t_{l-1}, t_{l+1}]$ for $l = 1, \dots, L$. We accept \mathbf{t}^* with probability

$$\min \left\{ \frac{L(\mathbf{t}^* | \mathbf{Y}_S, \mathbf{Y}_E, \mathbf{Y}_T, \boldsymbol{\beta}_S, \boldsymbol{\lambda}) (t_l^* - t_{l-1}) (t_{l+1} - t_l^*)}{L(\mathbf{t} | \mathbf{Y}_S, \mathbf{Y}_E, \mathbf{Y}_T, \boldsymbol{\beta}_S, \boldsymbol{\lambda}) (t_l - t_{l-1}) (t_{l+1} - t_l)}, 1 \right\}$$

5. Sample $\mathbf{t}, \boldsymbol{\lambda}, L | \mathbf{Y}_S, \mathbf{Y}_E, \mathbf{Y}_T, \mathbf{x}$ via a Metropolis-Hastings-Green move by proposing adding a split point and deleting a split point.

- **Birth Move:** Draw a random split point via a $t_* \sim U[0, t_{\max}]$ and set $\mathbf{t}^* = \text{Sort}(t_*, \mathbf{t})$. Changing the dimension of \mathbf{t} also adjusts the entries of $\boldsymbol{\lambda}$ in the following manner. Draw $U \sim U[0, 1]$ and assume that $t^* \in (t_{l-1}, t_l]$. Then we define the multiplicative perturbation as $\frac{\exp(\lambda_{l+1}^*)}{\exp(\lambda_l^*)} = \frac{1-U}{U}$ as in Green (1996). Then the new log heights of the hazard function are determined as

$$\lambda_l^* = \lambda_l - \frac{t_l - t_*}{t_l - t_{l-1}} \log \left(\frac{1-U}{U} \right)$$

$$\lambda_{l+1}^* = \lambda_l + \frac{t_* - t_{l-1}}{t_l - t_{l-1}} \log \left(\frac{1-U}{U} \right)$$

The proposed vectors $(\mathbf{t}^*, \boldsymbol{\lambda}^*)$ is accepted with probability

$$\frac{N(\lambda_{l+2}^* | \lambda_{l+1}^*, \sigma_\lambda^2) N(\lambda_{l+1}^* | \lambda_l^*, \sigma_\lambda^2) N(\lambda_l^* | \lambda_{l-1}^*, \sigma_\lambda^2) Poi(L+1 | \zeta_S) (2L+3)(2L+2)(t_* - t_{j-1})(t_j - t_*)}{N(\lambda_{l+1} | \lambda_l, \sigma_\lambda^2) N(\lambda_l | \lambda_{l-1}, \sigma_\lambda^2) Poi(L | \zeta_S) t_{\max}^2 U(1-U)(t_l - t_{l-1})}$$

times the likelihood ratio for $\mathbf{t}^*, \boldsymbol{\lambda}^*$ if it is < 1 , otherwise with probability = 1, from their previous values.

- **Death Move:** Similar to a Birth move, a death move adjusts both \mathbf{t} and $\boldsymbol{\lambda}$. We propose deleting one entry t_1, \dots, t_L with equal probabilities. Assume we delete t_l to obtain \mathbf{t}^* then we delete λ_{l+1} from $\boldsymbol{\lambda}^*$ and set

$$\lambda_l^* = \frac{\lambda_l(t_l - t_{l-1}) + \lambda_{l+1}(t_{l+1} - t_l)}{t_{l+1} - t_{l-1}}$$

We draw $U \sim U[0, 1]$ as the random perturbation to maintain balance between the two parameter spaces and we accept the proposed vectors $(\mathbf{t}^*, \boldsymbol{\lambda}^*)$ with probability

$$\frac{N(\lambda_{l+1} | \lambda_{l-1}, \sigma_\lambda^2) Poi(L-1 | \zeta_S) (t_{l+1} - t_{l-1}) t_{\max}^2 U(1-U)}{N(\lambda_{l+1} | \lambda_l, \sigma_\lambda^2) N(\lambda_l | \lambda_{l-1}, \sigma_\lambda^2) Poi(L | \zeta_S) (2L+1) 2J(t_{l+1} - t_l)(t_l - t_{l-1})}$$

times the likelihood ratio, if it is < 1 , otherwise with probability = 1,

C List of Phase I-II/III Design and Model Parameters

The design and model parameters needed to conduct a phase I-II/III trial are as follows:

- N_{ET} : The number of patients to enroll in the phase I-II portion of the trial.
- N^F : The number of patients assign optimized doses in phase I-II based on $\phi(\cdot)$ prior to adaptive randomization.
- $\phi(\cdot, \cdot)$: The desirability function for a given phase I-II data set. This is determined from the clinician elicited efficacy toxicity tradeoff contours. The clinician should specify a desired efficacy probability for a patient with a toxicity, a desired toxicity probability for a patient with efficacy and a pair of equally desirable efficacy and toxicity probabilities. $\phi(\cdot, \cdot)$ is determined based on a contour determined from the usual Eff-Tox methodology.
- ζ_{ET} : The hypermeans and standard deviations used in the Eff-Tox model (3), (4).
- $\underline{\pi}_E, \bar{\pi}_T, p_E, p_T$: The acceptable limits for toxicity and response probabilities in the context of continuing the trial. This also is used in the adaptive randomization scheme, only randomizing to acceptable doses.
- N_S : The maximum number of patients to enroll in the phase III portion of the trial.
- n_3^*, \dots, n_K^* : The respective numbers of patient events required to make interim looks for superiority and futility decisions at stages $k = 3, \dots, K$. These are determined based on desired information proportions. This portion of the trial should be planned considering the possibility that patient data will be lost after switching doses. We do not want the information proportion to be small so that a superiority/inferiority decision is not made based on a small number of patients treated at $A(\hat{x}_S^{\text{opt}})$.

- Superiority and Inferiority bounds $\bar{u}_k, \underline{u}_k$: The trial planner needs to determine the superiority and futility boundaries required to determine the outcome of the trial. This is done using East software.
- Δ : Desired improvement in mean survival to conclude $A(x_j)$ is better than C . We test the hypotheses $H_0 : \mu_C = \mu_{A(x_j)}$ versus $H_1 : |\mu_C - \mu_{A(x_j)}| > \Delta$.
- n_2^* : The number of patient events in the phase 3 portion of the trial to decide what dose $A(\hat{x}_S^{\text{opt}})$ should be used for the randomized comparisons to C . Clinicians should be willing to risk the possibility of losing patient data in the trial if a dose is discontinued where patients in phase 3 have been enrolled. This will likely reduce power from that had the correct dose of A been chosen initially.
- ζ_S : The prior mean number of split points in the baseline hazard. We suggest a default of $\zeta_S \in \{3, 4, 5, 6, 7\}$, since we have a lot of patient data to inform the posterior if no split points are needed.

Some of these values are set by the clinician to meet safety requirements and practical constraints, while others can be determined empirically via simulation. Functions in the package *Phase123* simulate repetitions of the Eff-Tox trial with adaptive randomization, re-optimize doses in phase III, and record final decisions about superiority and futility for both the conventional phase I-II \rightarrow III and phase I-II/III designs. The statistician should simulate trials under a reasonable set of different dose-efficacy and dose-toxicity probability configurations, relationships between (x, Y_E, T_T) and h_S , as well as survival time distributions.

D Guide for Using the *Phase123* Package

Here we describe in detail how to use several functions from the package *Phase123*. Explicit documentation of the arguments of each function and what they do can be found on CRAN. There are five important functions for the typical user, with some built upon functions that can be used for simulation studies on a cluster unit. These are:

1. **SimPhase123**: This function simulates replications of a phase I-II/III trial, given desired maximum sample sizes, assumed accrual rates, true efficacy and toxicity probabilities at each dose, time to event distribution family and true linear term coefficients related to dose, efficacy status and toxicity status. It returns a list containing two matrices containing results for both the phase I-II/III trial and the conventional phase I-II \rightarrow phase III paradigm. The first element of the list contains columns corresponding to \hat{x}_S^{opt} , the decision made in the trial (-1 indicates C is declared superior, 1 indicates A is declared superior and 0 indicates the trial ended in futility), the sample size of the phase III portion of the trial and the trial duration. The second element of the list contains the results for the conventional paradigm, with the first column instead containing $\hat{x}_{ET}^{\text{opt}}$.
2. **ReturnOCS**: This function takes output from the SimPhase123 function, true mean survival times for each dose, the true control mean, desired improvement in survival and the hypothesis (1 for alternative, 0 for null). If operating characteristics under the alternative hypothesis are desired, γ_1, γ_2 , and \bar{W} are returned. Otherwise, α is returned. Average sample sizes and trial durations also are returned.
3. **Reoptimize**: This function returns the newly reoptimized dose \hat{x}_S^{opt} as a function of the survival data $(Y^0, \delta), Y_E, Y_T$ and x_j for all patients treated with agent A during both phase I-II and phase III.
4. **AssignEffTox**: This function assigns a dose to the next cohort of patients in the phase I-II portion of the trial when patients are not adaptively randomized.
5. **RandomEffTox**: This function adaptively randomizes the next cohort of patients to a dose during phase I-II, with probabilities proportional to the desirability scores of the doses.

D.1 Simulating a Phase123 Trial and Evaluating Operating Characteristics

The following code comes from the example listed on CRAN for SimPhase123, with the number of simulation replications adjusted to $n\text{Sims} = 100$ for illustration. The function ReturnOCS is used to process and return the results. This example is scenario 3 in the manuscript for the exponential case under the alternative hypothesis.

```

##This is scenario 3 for the exponential case
##the additional phase 123 parameters and simulation parameters
# #####PHASE12 Parameters #####
  DoseStart=1
##True Efficacy and Toxicity Probabilities
PT = c(.05,.08,.1,.15,.2)
PE=c(.2,.25,.35,.4,.55)
##Dose Levels considered
Dose = c(1,2,3,3.5,5)
##Max Sample Size
NET=60
##Number of patients before randomization
NF=15
##Cohort size
cohort=3
##Hypermeans for Eff-Tox
Hypermeans = c(.022,3.45,0,-4.23,3.1,0)
Hypervars = c(2.6761, 2.6852, .2, 3.1304, 3.1165, 1)
Hypervars=Hypervars^2
##Contour Vector
Contour = c(.35, .75,.7,.4)
##Acceptability Criteria
PiLim = c(.3,.4)
ProbLim=c(.1,.1)
##Phase 12 accrual rate
Accrue12=5
####How long is the time window in phase 12?
Time12=1
#####PHASE3 Parameters#####
Nmax=500
##Number of patient events for interim looks
NLook = c(200,300,400)
##Superiority Boundaries
Sup = c(2.96, 2.53,1.99)
##Futility Boundaries (0 means no futility decision)
Fut = c(0,1.001,0)
##Average accrual rate for phase III
Accrue = 10
#####Phase123 Parameters#####
  ##Number of patient events to re-optimize doses
  NLookSwitch=50
# ##Time in between phase 12 and phase 3
  Twait=1
#####Simulation Parameters#####
##Family of Distributions
Family="Exponential"
##Shape parameter, Not needed for Exponential
alpha=1
##True Beta vector (beta_1,exp(beta_E),-exp(beta_T),beta_2,beta_0)
betaA = c(.1, .3, -1,-1,3.6)
##True beta vector for (exp(beta_E),-exp(beta_T),beta_C) of the control treatment
betaC=c(.3,-1,log(24/1.035111))
##True efficacy and toxicity probability for control group
ProbC = c(.3,.1)

```

```

###Number of simulations to run
nSims=100
##Run Simulations
Results=SimPhase123(DoseStart,Dose,PE,PT,Hypermeans,Hypervars,Contour,
PiLim,ProbLim,NET,NF,Accrue12,Time12,cohort,betaA,ProbC,betaC,
Family,alpha,Nmax,Accrue,Twait,NLookSwitch,NLook,Sup,Fut,nSims)

##Now Process Results for the Operating Characteristics using ReturnOCS
##True Mean Survival Time Vector for A##
Means=c(6.9, 24.7, 38, 33.1, 6.3)
##True Control Mean Survival Time ##
CMu=24
##Desired Improvement in Survival##
Delta=12
##Null=0, or Alternative = 1 Hypothesis ##
Hyp=1
ReturnOCS(Results,Means,CMu,Delta,Hyp)

```

During the simulations, for each replication, the dose \hat{x}_S^{opt} that is chosen is printed, saying for example: Assign Patients receiving A to dose 3. The last lines of code that use the ReturnOCS function output the operating characteristics. The mean survival vector can be determined using the function ReturnMeansAgent which has additional documentation on CRAN. Output from the ReturnOCS function is shown below:

```

> ReturnOCS(Results,Means,CMu,Delta,Hyp)
TRIAL TIMES
Conventional Design[1] 3.252829
Phase123 Design[1] 4.821758
Number of Patients
Conventional Design[1] 372.03
Phase123 Design[1] 457.2
Percentage of times best dose selected
Conventional Design[1] 0.1
Phase123 Design[1] 0.85
Generalized Power
Conventional Design [1] 0.09
Phase123 Design [1] 0.74
bar{W} values: Average true improvement in survival
Conventional Design [1] 1.729
Phase123 Design [1] 10.997

```

We see the results for the simulations for both the conventional paradigm and the phase I-II/III design displaying, in order, the average trial times (in years), the average number of patients treated, the proportion of times that the optimal dose was selected, the generalized power, and the value of \bar{W} . When the option Hyp=0 is chosen, the value of \bar{W} is not printed and the type I error is listed instead of the generalized power. When two or more doses have truly superior mean survival, both γ_1 and γ_2 are listed. When determining what value n_2^* should be used, we want to have both a high probability of selecting the best dose and a high generalized power. Values of n_2^* that are too large will result in low generalized power, but values of n_2^* that are too small will result in poor optimal dose selection probability.

D.2 Re-optimizing dose during phase III

After n_2^* patient deaths have occurred, we look at all the data for patients treated with agent A in both phase I-II and phase III to determine if we should switch doses. Below is code to implement this step, where $\hat{x}_{ET}^{\text{opt}} = 5$ and the truly optimal dose in terms of mean survival is $x_S^{\text{opt}} = 3$. The data provided was printed out running the SimPhase123 function with scenario 3 described above.

```

##This is data from a simulated phase123 trial where dose 5 has been brought to phase III
#but dose 3 is optimal in terms of mean survival
##Doses Given
Dose = c(1,2,3,3.5,5)
Dose=(Dose-mean(Dose))/sd(Dose)
#Survival and Follow up times
Y=c(2.49, 3.44, 0.90, 0.49, 0.13, 6.18, 4.39, 1.56, 4.16, 1.01,
9.76, 2.87,4.01,4.58,2.17,9.33,9.15,9.06,0.34,1.01,
6.57, 8.04, 7.97, 0.17,3.74, 7.83, 2.10, 5.28,6.87, 2.62,
6.53, 6.41, 1.10, 6.01, 5.94, 5.57, 5.27, 5.18,0.47,4.46,
0.16, 3.59, 2.31, 3.16, 3.98, 3.89, 3.75, 2.99, 3.41,3.33,
3.26, 3.22,3.12, 2.99, 0.40, 1.28, 1.52, 2.58, 2.52,2.42,
2.37, 1.38,1.11,1.06,1.02,0.98,0.93, 0.41, 0.17, 0.14,
5.96, 6.49, 5.42, 10.78, 23.37, 23.37, 23.07, 6.26, 23.07, 22.49,
21.35, 5.83, 22.02, 6.58, 22.02, 7.02, 4.13, 21.01, 0.05, 3.72,
20.72, 10.49, 20.36, 4.50, 11.40, 20.29, 20.29, 2.47, 2.59, 19.90,
8.39, 18.70, 3.38, 18.14, 9.70, 14.45, 10.23, 5.72, 3.73, 0.46,
11.19, 1.84, 15.16, 0.47, 4.13, 1.01, 3.02, 11.04, 15.04, 5.82,
9.30, 14.03, 14.03, 14.03, 5.27, 4.98, 13.52, 0.50, 1.99, 5.64)
##Censoring Indicators
I=c(1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,0,0,1,1,1,1,
0,0,1,1,0,1,1,0,1,0,0,1,0,0,0,0,0,1,0,1,1,
1,1,0,0,0,1,0,0,0,0,0,0,0,1,1,1,0,0,1,0,1,0,
0,0,0,0,1,0,0,1,1,1,1,0,0,0,1,0,0,1,1,0,1,
0,1,1,0,1,1,0,1,0,1,1,0,0,1,1,0,1,0,1,0,1,
1,1,1,1,1,1,1,1,1,1,1,1,1,1,0,1,1,0,0,0,1,1,
0,1,1,1 )
#Dose numbers given
Doses=c(5,5,5,5,5,5,5,5,5,5,5,5,5,5,5,5,5,5,5,5,5,5,5,5,5,5,
5,5,5,5,5,5,5,5,5,5,5,5,5,5,5,5,5,5,5,5,5,5,5,5,5,5,5,5,5,5,5,
5,5,5,5,5,5,1,1,1,2,2,2,3,3,3,4,4,4,4,4,4,4,4,4,2,2,2,3,3,3,4,4,4,3,
3,3,3,3,3,4,4,4,5,5,5,5,5,5,5,5,5,5,5,5,3,3,3,4,4,4,3,3 ,3,5,5,5)
##Efficacy Status Vector
YE=c(0,1,0,1,1,1,1,1,1,0,0,1,0,0,1,1,1,1,1,1,0,1,1,0,0,0,1,1,1,0,0,0,0,1,
1,0,1,1,0,1,0,0,0,1,1,1,0,0,0,1,0,0,0,1,0,0,0,1,0,0,0,1,1,1,
1,0,0,1,1,1,0,0,0,0,0,0,1,0,0,1,1,1,1,0,0,1,0,1,1,0,0,0,1,0,1,0,0,
1,0,0,1,1,0,0,1,0,1,1,0,1,1,0,1,0,1,0,0,1,1,0,0,0,1,1,0,0,0,1,0)
##Toxicity Vector
YT=c(1,0,0,0,0,0,0,0,1,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,
0,0,0,0,0,0,1,1,1,0,0,1,0,0,0,0,0,0,0,0,0,0,0,1,0,0,1,0,1,0,
0,0,0,0,0,1,0,0,0,0,0,0,0,0,0,0,0,0,1,0,1,0,0,1,0,0,1,0,0,0,0,0,0,0,0,0,0,0,0,0,0,
0,0,0,0,1,0,1,0,0,0,0,0,0,0,0,0,1,0,1,0,0,0,1,0,0,0,1,0,0,0,0,0,1,0,0,0,0,1,0)
##Hypermeans for Eff-Tox
Hypermeans = c(.022,3.45,0,-4.23,3.1,0)
Hypervars = c(2.6761, 2.6852, .2, 3.1304, 3.1165, 1)
Hypervars=Hypervars^2
###Number of iterations
B=2000
Reoptimize(Y,I,YE,YT, Doses, Dose, Hypermeans, Hypervars,B)

```

This function returns:

Assign Patients receiving A to dose [1] 3

The design accurately switched the dose of A for comparison to the control C from dose 5 to dose 3, that is $\hat{x}_S^{\text{opt}} = x_S^{\text{opt}} = 3$. After this switch, all patients randomized to A should be assigned dose 3 and those patients previously treated with dose level 5 should be excluded from the subsequent group sequential superiority and inferiority decisions.

D.3 Assigning patients during phase I-II

During phase I-II, the acting statistician assigns new cohorts of patients to a dose either deterministically or randomized with probabilities proportional to the acceptable desirability scores of the doses. This is also a function of the doses tried so far during phase I-II, since we do not allow escalating more than one untried dose at a time. While the number of patients enrolled in the trial is less than N^F we use the function `AssignEffTox` to assign new cohorts to a dose, which is also used to determine $\hat{x}_{ET}^{\text{opt}}$ after phase I-II (stage 1). After N^F patients are enrolled, we begin to use the function `RandomEffTox` to adaptively randomize patient cohorts to different doses. A very key point is that the dose numbers outputted by these functions are indexed starting at 0. This means that dose level 0 is really dose level 1 and dose 2 is really dose level 3, etc. This is because these two functions are also used in the c++ code for simulating phase I-II/III trials and c++ indexes vectors starting at 0. Below is the example code listed on CRAN along with output for both the `AssignEffTox` and `RandomEffTox` functions.

```
##Doses of patients currently enrolled
Doses= c(1,1,1,2,2,2,1,1,1,3,3,3,1,1,1,2,2,2)
##Efficacy indicators
YE = c(0,0,1,1,1,0,0,0,0,1,1,0,0,1,1,1,0)
##Toxicity Indicators
YT=c(0,0,0,1,1,0,1,0,0,1,1,1,0,0,0,1,0,0)
##Vector of Numerical Doses
Dose = c(1,2,3,3.5,5)
Dose=(Dose-mean(Dose))/sd(Dose)
##Five doses, but only 3 tried so we have
DosesTried=c(1,1,1,0,0)
## Contour Vector. (Prob Eff | No Tox, Prob Tox | Eff, PE pair, PT pair)
Contour = c(.35, .75, .7, .4)
##Hypermeans
Hypermeans = c(.022,3.45,0,-4.23,3.1,0)
Hypervars = c(2.6761, 2.6852, .2, 3.1304, 3.1165, 1)
Hypervars=Hypervars^2
##Acceptability Criteria
PiLim = c(.3,.4)
ProbLim=c(.1,.1)
##Number of iterations for MCMC
B=2000
AssignEffTox(YE,YT, Doses, Dose, DosesTried, Hypermeans, Hypervars, Contour, PiLim, ProbLim, B )

[1] 1

RandomEffTox(YE,YT, Doses, Dose, DosesTried, Hypermeans, Hypervars, Contour, PiLim, ProbLim, B )

[1] 0

##Note 0 means dose 1, 1 means dose 2, etc
```

According to the output of these two functions, if we were deterministically assigning doses (i.e. $n < N^F$) we would assign the next cohort to dose 2. If we randomly assigned this cohort with probabilities proportional to the desirability scores, we would assign the next cohort to dose level 1.

E Web Tables

Table 1: Simulation results for the phase I-II Eff-Tox portion of phase I-II/III, for $N_{ET} = 90$. True outcome probabilities $(\pi_E, \pi_T)^{TR}$, desirabilities $\phi\{(\pi_E, \pi_T)^{TR}\}$, and operating characteristics for the usual (non-adaptively randomized) EffTox phase I-II trial design.

Scenario	Value	1	2	3	4	5
1	$(\pi_E, \pi_T)^{TR}$	(.20, .10)	(.40, .15)	(.60, .25)	(.65, .35)	(.70, .50)
	$\phi\{(\pi_E, \pi_T)^{TR}\}$	-.37	-.13	.05	-.01	-.13
	% Selected	.01	.14	.31	.31	.24
	# Treated	7.9	19.3	26.6	23.8	12.3
2	$(\pi_E, \pi_T)^{TR}$	(.2, .05)	(.25, .08)	(.35, .10)	(.40, .15)	(.55, .20)
	$\phi\{(\pi_E, \pi_T)^{TR}\}$	-.30	-.26	-.14	-.13	.04
	% Selected	.02	.04	.07	.11	.76
	# Treated	7.8	11.1	18.4	21.6	31.1
3	$(\pi_E, \pi_T)^{TR}$	(.40, .10)	(.50, .15)	(.60, .35)	(.65, .60)	(.70, .70)
	$\phi\{(\pi_E, \pi_T)^{TR}\}$	-.06	.03	-.09	-.35	-.40
	% Selected	.34	.50	.12	.03	.01
	# Treated	22.9	33.1	21.4	10.7	1.9

Table 2: Simulation results for phase I-II sample size $N_{ET} = 90$ in the phase I-II/III design. Values are repeated for $N_{ET} = 60$ from Table 5 in the manuscript to facilitate comparison. α is the probability of a type I error or concluding an inferior version of A is better than C under the null. γ_1 is the probability of the generalized power event at $A(x_S^{opt})$ (selecting the best dose x_S^{opt} and declaring it to be superior to C) under the alternative hypothesis. γ_2 is the generalized power (probability of selecting any truly superior dose of A and declaring it superior to C). \bar{W} is the mean improvement in patient survival time under the alternative hypothesis, \overline{Dur} is the mean trial duration, and \bar{N} is the mean sample size.

Scenario	N_{ET}	Alternative Hypothesis					Null Hypothesis			
		\bar{W}	γ_1	γ_2	\overline{Dur}	\bar{N}	α	\overline{Dur}	\bar{N}	
1	60	10.15	.83	.83	4.73	479.2	–	.03	4.32	492.0
	90	10.81	.88	.88	4.71	478.2	–	.03	4.34	493.7
2	60	8.97	.75	.75	4.45	470.7	–	.02	4.18	489.9
	90	10.02	.84	.84	4.46	469.0	–	.02	4.16	489.4
3	60	11.51	.79	.79	4.56	476.9	–	.04	4.22	485.6
	90	12.66	.88	.88	4.56	476.8	–	< .01	3.57	422.2
4	60	5.86	.42	.42	4.30	472.0	–	.05	3.81	442.0
	90	6.51	.46	.46	4.16	464.4	–	.05	3.93	456.9
5	60	16.71	.68	.88	4.24	464.4	–	.03	4.37	493.9
	90	18.25	.81	.94	4.15	462.1	–	.03	4.37	493.8
6	60	12.67	.59	.75	4.53	472.7	–	.04	4.39	494.0
	90	13.99	.67	.90	4.46	471.1	–	.04	4.41	494.8

Table 3: Robustness Study Parameters: True means under null and alternative hypotheses for each distribution considered. $\boldsymbol{\mu}_{H_m} = (\mu_{A(x_1)}, \mu_{A(x_2)}, \mu_{A(x_3)}, \mu_{A(x_4)}, \mu_{A(x_5)})$ are the mean survival times for doses x_1, \dots, x_5 of A under hypothesis $m = 0, 1$.

Scenario	Distribution	$\boldsymbol{\mu}_{H_0}$	$\boldsymbol{\mu}_{H_1}$
1	Exponential	(8.3, 17.9, 24, 22.5, 9.8)	(1, 14.5, 36.2, 28.3, 1)
	Lognormal, $\sigma = .5$	(9.3, 17.9, 24.0, 22.5, 9.8)	(1.1, 16.4, 41.0, 32.1, 1.2)
	Lognormal, $\sigma = 1$	(8.4, 17.9, 24.0, 22.5, 9.8)	(1, 14.7, 36.7, 28.9, 1.0)
	Weibull Increasing	(8.3, 17.9, 24.0, 22.5, 9.8)	(1.1, 15.4, 38.5, 30.2, 1.1)
	Weibull Decreasing	(8.3, 17.9, 24.0, 22.5, 9.8)	(1.0, 14.7, 36.8, 28.9, 1.0)
	Gamma	(8.3, 17.9, 24.0, 22.5, 9.8)	(1.0, 14.7, 36.8, 28.9, 1.2)
2	Exponential	(3.2, 10.1, 20.4, 24, 20.4)	(3.0, 12.9, 31.8, 40, 36.6)
	Lognormal, $\sigma = .5$	(3.2, 10.1, 20.4, 24.0, 20.4)	(3.2, 13.7, 33.9, 42.6, 38.9)
	Lognormal, $\sigma = 1$	(3.2, 10.1, 20.5, 24.0, 20.5)	(3.1, 13.0, 32.0, 40.3, 36.8)
	Weibull Increasing	(3.2, 10.1, 20.4, 24.0, 20.5)	(3.1, 13.0, 32.1, 40.4, 36.8)
	Weibull Decreasing	(3.2, 10.1, 20.4, 24.0, 20.4)	(3.1, 13.0, 32.1, 40.4, 36.9)
	Gamma	(3.2, 10.1, 20.4, 24.0, 20.4)	(3.1, 13.3, 32.8, 41.2, 37.7)
3	Exponential	(14.0, 17.8, 21.9, 23, 24)	(7.1, 10.3, 16.0, 19.5, 36)
	Lognormal, $\sigma = .5$	(14.1, 17.8, 21.9, 23.0, 24.0)	(8.0, 11.7, 18.1, 22.0, 40.8)
	Lognormal, $\sigma = 1$	(14.1, 17.8, 21.9, 23.0, 24.0)	(7.2, 10.6, 16.3, 19.9, 36.8)
	Weibull Increasing	(14.0, 17.8, 21.8, 22.9, 24.0)	(7.2, 10.5, 16.2, 19.7, 36.5)
	Weibull Decreasing	(14.0, 17.8, 21.9, 23.0, 24.0)	(7.2, 10.5, 16.2, 19.8, 36.6)
	Gamma	(14.0, 17.8, 21.9, 23.0, 24.0)	(7.2, 10.5, 16.2, 19.8, 36.6)
4	Exponential	(9.5, 18.5, 24, 22.5, 10.4)	(6.9, 24.7, 38, 33.1, 6.3)
	Lognormal, $\sigma = .5$	(9.5, 10.6, 24.0, 22.6, 10.4)	(7.0, 25.2, 38.8, 33.8, 6.5)
	Lognormal, $\sigma = 1$	(9.5, 18.6, 24.0, 22.6, 10.4)	(7.0, 25.0, 38.6, 33.6, 6.4)
	Weibull Increasing	(9.5, 18.5, 24.0, 22.5, 10.4)	(7.0, 25.0, 38.6, 33.7, 6.4)
	Weibull Decreasing	(9.5, 18.5, 24.0, 22.6, 10.4)	(6.9, 24.9, 38.3, 33.4, 6.4)
	Gamma	(9.5, 18.5, 24.0, 22.7, 10.4)	(6.9, 24.9, 38.3, 33.4, 6.4)
5	Exponential	(9.3, 18.7, 24.0, 22.7, 10.4)	(7.8, 28.8, 44, 38.6, 7.4)
	Lognormal, $\sigma = .5$	(9.3, 18.7, 24.0, 22.7, 10.4)	(7.8, 28.8, 44, 38.6, 7.4)
	Lognormal, $\sigma = 1$	(9.3, 18.7, 24.0, 22.7, 10.5)	(7.8, 28.9, 44.2, 38.8, 7.4)
	Weibull Increasing	(9.3, 18.7, 24.0, 22.7, 10.5)	(8.0, 29.6, 45.2, 39.7, 7.6)
	Weibull Decreasing	(9.3, 18.7, 24.0, 22.8, 10.5)	(7.9, 29.3, 44.8, 39.3, 7.5)
	Gamma	(9.3, 18.7, 24.0, 22.8, 10.5)	(7.9, 29.3, 44.8, 39.3, 7.5)
6	Exponential	(24, 13.6, 8.9, 6.8, 7.8)	(38, 24.6, 18.4, 15.0, 21.1)
	Lognormal, $\sigma = .5$	(24.0, 13.6, 8.9, 6.8, 7.8)	(38.0, 24.5, 18.4, 15.0, 21.0)
	Lognormal, $\sigma = 1$	(24.0, 13.6, 8.9, 6.8, 7.8)	(38.9, 25.2, 18.8, 15.4, 21.6)
	Weibull Increasing	(24.0, 13.6, 8.9, 6.8, 7.9)	(38.2, 24.7, 18.5, 15.1, 21.2)
	Weibull Decreasing	(24.0, 13.6, 8.9, 6.8, 7.9)	(38.7, 25.0, 18.7, 15.3, 21.4)
	Gamma	(24.0, 13.6, 8.9, 6.8, 7.9)	(38.7, 25.0, 18.7, 15.3, 21.4)